

UNCLASSIFIED

Defense Technical Information Center Compilation Part Notice

ADP014135

TITLE: Risk Assessment Methodologies for Fracture-Critical Components

DISTRIBUTION: Approved for public release, distribution unlimited

Availability: Hard copy only.

This paper is part of the following report:

TITLE: Aging Mechanisms and Control. Symposium Part A -
Developments in Computational Aero- and Hydro-Acoustics. Symposium
Part B - Monitoring and Management of Gas Turbine Fleets for Extended
Life and Reduced Costs [Les mecanismes vieillissants et le controle]
[Symposium Partie A - Developpements dans le domaine de
l'aeroacoustique et l'hydroacoustique numeriques] [Symposium Partie B ...

To order the complete compilation report, use: ADA415749

The component part is provided here to allow users access to individually authored sections of proceedings, annals, symposia, etc. However, the component should be considered within the context of the overall compilation report and not as a stand-alone technical report.

The following component part numbers comprise the compilation report:

ADP014092 thru ADP014141

UNCLASSIFIED

Risk Assessment Methodologies for Fracture-Critical Components

A.D. Boyd-Lee and D.P. Shepherd

Structures and Materials Centre, QinetiQ, Farnborough, Hants, GU14, 0LX, UK

Phone +44-1252-392000

Fax +44-1252-397298

Email: adblee@qinetiq.com, dpshepherd@qinetiq.com

1 INTRODUCTION

Within gas-turbine aeroengines, fracture-critical components are defined as those whose in-service failure would hazard the entire aircraft. For these components, airworthiness regulations require that a maximum permitted service life be identified, such that the probability of failure occurring before this life is reached is extremely remote. However, although the intent behind the derivation of these lives is to minimise the possibility of in-service failure, the procedures used to establish them are not usually specified in terms of this probability of dysfunction itself. For example, under the conventional 'safe life' methodology, it is the probability of a component exceeding some identifiable fraction of its total life, rather than of the total life, which is used as a fundamental criterion. The portion of the life remaining beyond this point is regarded as an additional, unspecified, safety factor, which ensures the probability of actual failure is acceptably small. Similar considerations apply to most other methods of component life determination.

There are, however, many instances that can arise during the service life of an engine fleet, in which it becomes very desirable to ascertain the actual probability of a service component failure. This is often prompted by the occurrence of some unforeseen or unpredictable event, which indicates that the actual failure probabilities are higher than those that should be achieved under the normal airworthiness procedures. In such circumstances, informed decisions concerning the appropriate action can only be made if accurate estimates of the resulting failure probability (usually referred to as the risk) are estimated. For example, if the safe life for a certain component is cut during the service life of an engine (which can happen for a variety of reasons) a situation may result where over life components are being operated in service. Since the immediate rejection of these components may present severe operational difficulties, the question as to how fast they must be removed whilst maintaining acceptable safety levels becomes extremely important. Similarly, very occasionally, components belonging to a particular set or batch are discovered subsequent to entry into service to be substandard. Again, the affected components should be dealt with as quickly as possible, but how this process should be managed depends heavily on how seriously they impact on safety. For these and other cases of risk exposure, the required action can only be properly assessed if the reduction in the level of safety is estimated in terms of effective measures of the probability of actual failure over specified time periods.

To illustrate, some of the considerations associated with estimation of probabilities of failure, the next section discusses several scenarios associated with engine service operation for which such assessments become desirable. Aspects commonly encountered in deriving the appropriate estimates are described, and solution methods are discussed briefly. These are illustrated by examples of actual risk assessment exercises, which have been conducted in support of UK military engine operation. Subsequent sections analyse how judgements of the significance of airworthiness risks are made. Particularly, the manner in which the risks are to be related to the RAF Hazard Risk Index (HRI) are discussed. It is emphasised that these may vary significantly depending on the situation being addressed. Finally, these issues are related to UK current and future safety requirements at the aircraft platform level and the implications discussed.

2 RISK ASSESSMENT METHODOLOGIES

There are diverse factors that can cause the risk of engine failure to exceed normal levels, thus requiring the explicit calculation of failure probabilities to become necessary. These include but are not limited to: foreign object damage (FOD), problems arising in the critical parts lifing process itself, problems

associated with production and manufacturing process, errors in maintenance, hazards associated with extreme operating conditions. From the viewpoint of the methods needed to assess the risks that arise, these factors fall into a number of different categories. At the simplest level, there are hazards that are ever present, and unaffected by parameters associated with the engine itself. Examples of this type may include those attributable to FOD (foreign object damage) or extreme weather conditions. The calculations for the associated risks can be simple in such cases, for example, obtained by dividing the number of occurrences across the fleet by an appropriate measure of exposure time (usually engine flying hours). However, for other cases, the calculation procedure is more complex. For example, to address potential failure modes associated with life-limiting mechanisms, account must be taken of the dependence of the risk of failure upon the life consumed, which in turn could be dependent on a variety of other factors such as component tolerances and the severity of service usage. Moreover, in situations where some subset of the population of components is below standard in some sense, difficulties arise in identifying and characterising the degree of non-conformance. This also occurs in problems associated with a lapse in maintenance standards, or some difficulty with build, or the use of contaminated fuels or lubricants. Finally, problems associated with operational factors may require special methods. Each of these latter cases will be considered in more detail below.

In any risk assessment associated with fracture-critical parts, the method used in deriving the risk estimates will depend significantly on the procedure used in deriving the original component service lives. This is because the form of the available data for calculating the estimates will be determined to a large extent by the particular lifing methodology employed. In this paper, discussion of problems associated directly with components will concern those lifed according to the 'safe life' methodology, whereby lives are directly inferred from suitable test data using representative statistical models. In these situations, the statistical models required for the risk estimation procedure can be derived from those which underlie the procedures themselves. The reason for this focus on 'safe life' procedures is that, for the most part, risk assessment exercises are required for engine fleets which have reached some degree of maturity, and it is only relatively recently that alternative methodologies have been incorporated within UK airworthiness regulations. However, it is important to note that risk calculation exercises conducted for parts lifed using alternative methods may require quite different models and methods.

2.1 Risk assessment associated with Low Cycle Fatigue life shortfall

It is a normal part of aeroengine certification procedures that acceptable safe service lives for the fracture-critical components are identified prior to engine entry into service. Particularly for the larger fleets, it is highly likely that some lives will be revised during the operational life of the fleet. For UK military engines in particular, such a life progression is normally a required part of the life management procedure. This is in recognition of the fact that the true nature of component degradation can be better quantified and understood by taking account of service experience. Thus, it has been the case that when a new aeroengine design enters service, the service lives declared for its fracture-critical components have been set at half of the respective 'safe lives' calculated from the component rig tests. The remainder is only released on completion of additional tests of components that have been exposed to significant service usage. Moreover, there are several other factors not specified by the airworthiness procedures, which ensure that component safe service lives are continuously evolving as the engine progresses through its life cycle. For example, the stress analysis techniques on which all life declaration is based are themselves subject to an ongoing process of development, and it is common practice for these stress analyses to be recalculated. Indeed, when improved methods have been applied to service components, very large life revisions have sometimes resulted. Also, the calculation of the exchange rates, used to convert from engine flying hours to cycles, are periodically recalculated as the available data sets of sampled flights increase over time.

Since life revisions occur more often when the fleet is at a mature age, a sufficiently large downward revision can cause some of the respective service components to become immediately life expired. In such situations, the task of managing the withdrawal of these components is a question of minimising the operational impact, whilst keeping the risks below acceptable limits. This can only properly be undertaken if the risks themselves have been estimated using the most effective available methods.

In order to perform such a risk assessment, the fundamental task (as in most situations) is to identify the relevant failure distribution. For the case of straight low cycle fatigue life shortfall, this process is assisted by the fact that a standard statistical model for describing the life distribution of fracture-critical

parts already exists. That is the model from which the lifing procedures given in current UK airworthiness regulations are derived. These procedures specify that the distribution of life to crack initiation (defined as the appearance of a crack of 0.75mm surface length) is lognormal, with known log standard deviation¹. All that is required to fully specify the model is to identify the log mean, which is obtained from appropriate test data. There is then the problem of identification of the dysfunction distribution. It is thus necessary to obtain the ratio of the life to dysfunction versus the life to crack initiation (for typical components the ratio is equal to 1.5). If it is further assumed that the dysfunction distribution is itself lognormal with the same given log standard deviation, it follows that the failure log mean is equal to this ratio times the known initiation log mean. This allows the dysfunction distribution to be completely specified, thus giving the probability of failure for any particular component lifetime.

Whilst these properties form the basis of most risk models describing low-cycle fatigue situations, there are a number of complicating features which must necessarily be addressed if meaningful risk estimates are to be derived. Detailed descriptions of the models which result from these considerations have been given elsewhere (ref.1), so the paper will only give a brief outline of some of the problems which arise and how they are addressed. Firstly, it has often been observed that the variability in crack initiation data is greater than that in crack propagation. This suggests that the assumption of equal variance between the initiation and failure distributions is unrealistic, and that in practice the variance on the failure distribution will be smaller than that specified for initiation. Consequently, an appropriate means of calculating the failure variance, given the variabilities in initiation and propagation, must be derived.

Particularly for fixed wing military aircraft and helicopters, minor cycles have a much greater relative effect on the rate of damage accumulation once a crack has initiated. Consequently, the exchange rates (that convert between engine flying hours and reference cycles) are several times larger during the crack propagation phase of component life than during crack initiation. This impacts on lifing models in two ways. Firstly, it means that, for all significant applications, the model is best expressed specifically in terms of hours (missions) rather than cycles. The reason for this is that a component lifetime is generally expressed in terms of hours, and the aim of the model is to evaluate the risk of failure. Further, suppose that the lifetime is large enough to have reached some significant percentile on the initiation distribution. If we want to convert this lifetime into cycles, we are forced to consider the possibility that the component might already have an initiated crack. Moreover, since the exchange rates for initiation and propagation are different, it is necessary to address when this event occurs in order to calculate the equivalent cycles. However, since the life to initiation is itself a random quantity, we have no idea when this will have occurred for a particular component. Consequently, we have no information, for a specific component, about what the equivalent lifetime in reference cycles is once the life gets large enough. The second impact of this inequality of exchange rates is that the relationship between the log means for the initiation and failure distributions will be very different, depending on whether they are specified in terms of cycles or hours. Indeed, the ratio between the means when considered in hours is very much less than 1.5; typically, it will be between 1.15 and 1.2 depending on the exact ratio of the exchange rates. A significant implication is that, the risk of failure rises very much more rapidly when considered in terms of actual flight hours than it appears to when expressed in terms of reference cycles.

An example of a more complex situation is one where fatigue cracks are initiated by corrosion. In terms of probability distributions, the degradation processes associated with corrosion differ from that of conventional low-cycle fatigue crack initiation. Whilst a component susceptible to corrosion may sustain numerous corrosion pits, the vast majority of these will not usually result in a propagating crack. However, occasionally, corrosion can continue beneath the surface, leading to the accelerated initiation of a large crack. Moreover, it is almost impossible to detect, from surface inspection, which of the corrosion features will result in this type of behaviour. In this situation, the conventional fatigue model of initiation and propagation is inappropriate, and if an accurate assessment of the worst case is used as the basis of the calculation, the life obtained will be unrealistically short. To model this situation, it is necessary to take account of both of the distribution in properties of the initiation mechanism, its probability of occurrence and its location. Thus, the calculated safe life will contain an expression describing the minimum life resulting from this mechanism, multiplied by an expression describing the probability that

¹ The log standard deviation of a lognormal distribution is the standard deviation of the associated normal distribution, obtained by taking the logarithm of the component lifetime. Similarly, the log mean is the mean of this distribution.

it will occur at all. Only by combining these elements can a realistic estimate of the overall risk be obtained.

2.2 Risk assessment for batch problems

A risk management situation that requires a different approach can be illustrated by an investigation prompted by the discovery of service components that are either inferior or substandard in some way. Manufacture of fracture-critical components involves detailed checking of the dimensions and surface profiles of new major components. More costly inspection techniques, such as microanalysis, X-ray analysis, SEM etc. are usually limited to batch sampling. Also, manufacturing quality control of aeroengine fasteners is usually based on sample inspection of batches. Hence, the possibility of a manufacturing basic inspection error should be extremely rare by virtue of the many pairs of quality controller hands that aeroengine components must go through. When a poor quality component has entered service, in most cases there appears to have been a factor that made the defect difficult to detect, such as:

1. very small inner radii (e.g. Helical spline case discussed below);
2. inaccessibility (e.g. threads, hidden interfaces, interior geometry, etc.);
3. critical dimensions (e.g. those that affect running clearances or component assembly); or
4. manufacture over 20 years ago when quality control was less stringent; or
5. inconsistencies that have been associated with the materials processing (e.g. large grains with twins in Ni-base superalloys, Waspaloy inclusions, soft alpha in Ti, plucking, etc).

Should such a feature or processing complication affect the performance of the component, the consequences can sometimes be such that high levels of risk result. An example of such a case was encountered in connection with the inner radii (at either side of the bottom of the teeth) of a helical spline on a shaft within a military aeroengine. Theoretical predictions and testing confirm that the expected life of this spline depends critically on these inner radii. Consequently, the design specification is that the inner radii should not be less than 0.3mm. Unfortunately, one of these splines was found in service having teeth with significantly smaller inner radii. The primary lifing issue was to estimate how the radius affected the life of the component. From this, the risk of shaft failure could be assessed and it was established that the minimum acceptable inner radius was in region of 0.15mm. This provided a safe criterion for acceptance or rejection of components.

It was also necessary to characterise the extent of the deviance within the population, so that the severity of the problem could be estimated. Characterisation of the distribution of radius sizes, in conjunction with the life estimation calculations, ensured that appropriate risk estimates were derived, and the offending batches of splines were identified and removed from service.

A second example of problems associated with production batches arises in connection with material that had been subject to inadequate process control. In one case, several cracks and materials defects were found in Waspaloy service discs and shafts manufactured about 30 years ago. The cracks were found to have initiated from extremely large inclusions that were the result of substandard materials processing. Since identification of the higher risk components could not be achieved by inspection, a process modelling approach was used instead. The plastic-flow forging process for the disc was predicted to identify where the inclusion had come from in the original casting. It was established that the oxide had fallen into the melt from the wall of the crucible. Having established this, the process model was then used to predict that the defects were more likely to be present in the rim and outer web than in the bore of the disc, which was consistent with the observations. A conservative probability of this event occurring was estimated, and from this the risk of in-service failure as a function of component life was evaluated.

In general, risk assessments involving batches of substandard components tend to be highly case dependent and the more complex cases demand a wide range of expertise to reach an acceptable solution. However, in any situation, it will almost certainly be necessary to characterise the severity of the problem. From this, the impact on the component lives can be estimated, and then the associated risks can be derived. Also, a significant issue can be deciding what, if any, inspections are needed to ensure that the problem is adequately characterised.

2.3 Risk assessment connected with operational usage

When a hazard associated with operational usage is in some sense exceptional, a risk analysis may be required. Typically, this will involve an interaction between an existing, identified hazard, and a particular operational situation or manoeuvre. An example is given by a power surge problem associated with a certain aeroengine. It was found that a voltage spike in the fuel system could result in a sudden uncontrollable rise in thrust. The probability of occurrence was remote (of the order of 10^{-5} per engine flying hour) and in most stages of flight it should not result in an air accident. However, if it were to occur during manoeuvres with little margin for error, it could result in the loss of the aircraft. For this problem, the available evidence indicated that occurrence of surge was independent of any normally measurable parameter (such as altitude, age of engine, etc) and in particular of the type of flying being undertaken. Thus, the estimated probability of aircraft loss per aircraft flying hour is the multiple of the probability of surge and the proportion of time spent in the critical manoeuvres. Simplifying slightly, the cumulative risk of a catastrophe was thus obtained by multiplying this probability by the time taken to rectify the fuel circuit problem.

The critical elements in this situation are the assessment of the operational usage, and the assumption of 'randomness' of the underlying event. The latter factor is particularly important, because deviations from this assumption could radically alter the predictions. Thus, in the above example, if it transpired that the fault could only arise under a particular condition or combination of conditions, then the underlying probability of aircraft loss could be zero if this could not occur in conjunction with the critical modes of flight. Alternatively, if it transpired that surge is more commonly associated with the critical flight modes then risk of catastrophe would be many times higher.

3 RISK MANAGEMENT USING THE RAF HAZARD RISK INDEX

3.1 Background

Once a particular risk has been estimated or updated, the next consideration is to assess the acceptability of this risk with respect to prioritisation of maintenance or technical actions.

For UK military aeroengines, the RAF Hazard Risk Index (HRI) controls the management of technical risks (ref. 2). The current HRI procedure is based around two matrices: one for fighter aircraft (figure 1) and the other for military transport or passenger aircraft (figure 2).

The primary purpose of the HRI is to ensure common standards of safety are applied across the RAF fleets. It is also used to prioritise technical actions, given limited resources. The associated HRI probability criteria have actively driven current risks of failure for at least the last 5 years, and consequently they have become a minimal requirement to support current RAF safety levels.

	PROBABILITY PER AIRCRAFT FLIGHT HOUR				
	(A) FREQUENT $>10^{-3}$	(B) PROBABLE 10^{-3} to 10^{-4}	(C) OCCASIONAL 10^{-4} to 10^{-5}	(D) REMOTE 10^{-5} to 10^{-6}	(E) IMPROBABLE $<10^{-6}$
(1) CATASTROPHIC Death or system loss				8	12
(2) CRITICAL Severe injury, severe occupational illness, or major system damage			6	10	15
(3) MARGINAL Minor injury, minor occupational illness, or minor system damage		9	11	14	17
(4) NEGLIGIBLE Less than minor injury, occupational illness or system damage	13	16	18	19	20




CRITERIA	
	UNACCEPTABLE
	UNDESIRABLE
	ACCEPTABLE WITHOUT REVIEW

Figure 1: RAF Hazard Risk Index – Combat aircraft

	PROBABILITY PER AIRCRAFT FLIGHT HOUR				
	(A) FREQUENT >10 ⁻³	(B) PROBABLE 10 ⁻³ to 10 ⁻⁵	(C) OCCASIONAL 10 ⁻⁵ to 10 ⁻⁷	(D) REMOTE 10 ⁻⁷ to 10 ⁻⁹	(E) IMPROBABLE < 10 ⁻⁹
(1) CATASTROPHIC Death or system loss				8	12
(2) CRITICAL Jeopardised flight, severe injury, major system damage				10	15
(3) MARGINAL Minor injury, minor system damage, or precautionary mission abort	7	9	11	14	17
(4) NEGLIGIBLE Mission safety unaffected	13	16	18	19	20




CRITERIA
 UNACCEPTABLE
 UNDESIRABLE
 ACCEPTABLE WITHOUT REVIEW

Figure 2: RAF Hazard Risk Index – Military transport aircraft

Each significant failure mode or other hazard is assigned to an HRI category. Red categories risks are deemed unacceptable, the amber ones require management action (undesirable) and the green ones require no action. The 'amber' band is there to provide a degree of flexibility in the management, and avoid unnecessary groundings and engine removals. A management objective is to reduce the risk associated with 'amber/red HRI category' failure modes to the extent that they become 'green HRI category' failure modes.

An aeroengine has numerous potential failure modes, and if all of these were allowed to exceed a risk of 10⁻⁶/afh, (i.e. per aircraft flight hour) the aircraft would be unsafe to fly. Thus, design standards such as Def. Stan. 00-971 specify a requirement for even lower probabilities of failure than those allowed by the HRI matrix. Typically, the effect is that only a few dozen of the failure modes of a particular engine type belong to the amber HRI category.

3.2 Interpretation of the current HRI matrix

3.2.1 Measures of risk

As originally defined, the HRI matrices have axes of fleet average risk per aircraft flying hour versus severity of the hazard. The fleet average is calculated simply by summing the risks over an entire engine fleet, and dividing by the number of aircraft. Many past life management exercises have been based upon this measure of risk.

In situations where a given risk is constant, (in other words, unaffected by other parameters associated with the operation of the engines such as component life or flight envelope), fleet average is an adequate measure of the significance of the risk. It is worth noting that, under these circumstances, the fleet average risk is the same as the maximum risk per aircraft flying hour in the fleet. When the severity of the risk is low, and the rate of arising is high, the average risk is particularly useful to assess trends in the fleet.

However, a major problem with using this risk measure as the default to decide levels of management action is that airworthiness risks are more often than not highly non-uniformly distributed. As a result, in some cases, the seriousness of a given situation is not indicated by the fleet average risk. To illustrate this, consider a situation in which the life of the fleet-leading component (of a given type) is just 15% higher than the declared safe life. Then the associated risks would typically be an order of magnitude higher than that at the declared life, and the vast majority of the risk could be confined to a life-leading component in a service aeroengine. Suppose further that the fleet average risk /afh is marginally below 10⁻⁶/afh and there are 500 engines in service. Assuming that 90% of the risk is being incurred by the life leading

component, then it would be incurring a risk of $4.5 \times 10^{-4}/\text{afh}$ ($= 0.9 \times 500 \times 10^{-6}$), that is HRI category 2 and clearly unacceptable.

In view of such considerations, it is suggested that no single measure of hazard exposure is appropriate to all situations, and in practice different measures are relevant to different circumstances. Moreover, sometimes multiple measures of risk are required to properly quantify a life management situation. In what follows, two alternative probability measures of incident arisings are described, and their use in various situations discussed.

3.2.2 Peak risk

The peak risk per aircraft flying hour is defined as the peak risk with respect to the parameters that cause a risk in question to be non-uniformly distributed. For example, in situations involving over life components, the peak risk will be the maximum with respect to the different components (i.e. the highest percentage over life component). Alternatively, if a risk were driven with respect to a particular operational manoeuvre or sortie type, then the peak risk would be the maximum amongst the risks for the different activities. Peak risks are important in situations where the risks (or the vast majority of the risks) are confined to a few components or engines, or a limited set of circumstances. In these situations, it is the peak risk that should be used in conjunction with the HRI, and not the fleet average risk. The advantage of using the peak risk in this way is that it clearly highlights the source of the problem, and will direct the management action to address those risks, which constitute the majority of the total. This ensures that no service component incurs unacceptable risk.

3.2.3 Expected cumulative arisings

The expected cumulative arisings per annum (or other relevant operational period) is defined as the cumulative risk over that period, summed over the entire fleet. When the risk is more evenly distributed, the expected cumulative arisings may be a more useful measure than peak risk, for two main reasons. Firstly, it takes account of the duration of risk exposure (usually associated with the time required to implement a maintenance action). This is very important when the risks under investigation are changing with time, and it allows for an assessment of how quickly problems need to be addressed. Secondly, it takes account of the number of aircraft affected. This is important in assessing priorities when different numbers of aircraft are affected by different problems. The expected cumulative arisings has the additional advantage that it is the measure that can most accurately related to observed failure rates. This is because it takes into account all aspects of the service situation including fleet size, observation period, engine mark and so on².

To illustrate the use of these risk measures, consider again the example discussed on page 3, whereby a component was subject to a life cut which resulted in a number of installed parts having exceeded their safe life. In this example, the peak risk for the fleet leading component was calculated, which, when compared to the HRI matrix resulted in an amber reading. Consequently, a programme of component withdrawals was derived, which required the components to be removed over a period of 2 years. However, it was also recognised that the risks associated with the installed components would be increasing whilst they were still in service, and it was important to ensure that these risks would not become unacceptable during this period. For this reason, an upper bound on the expected cumulative arisings over the period of the management plan was introduced, and the risk curves were integrated to ensure that the management plan conformed to this criterion.

4 TOWARDS AN IMPROVED HRI MATRIX

Given that several different measures of risk are used in practice to make appropriate aeroengine fleet managerial decisions, users of the HRI matrix have had difficulty interpreting airworthiness risk in relation to this decision matrix. Furthermore, difficulties of interpretation have also been encountered with the severity axis of the HRI matrix. It is suggested that the HRI matrix could be improved by using stronger inputs. These issues are discussed in turn.

²Note that for small fleets, the resources might not be available to revise component life and exchange rate calculations as frequently as for larger fleets. Thus, the uncertainty in estimates of the peak risk may be significantly higher, and so it may be necessary to introduce additional safety margins to account for this.

4.1 Risk criteria

Of the three measures of risk discussed, the peak risk and the expected cumulative arisings are the most critical with respect to safety. A solution would therefore be to use a dual risk axis for the HRI columns (figure 3). One axis gives the risk criteria in terms of peak risk (see section 3.2.2), the other in terms of expected cumulative arisings (section 3.3.3). The user would then select the lowest letter column associated with either risk. For example, if for a certain hazard the estimated peak risk is $2 \times 10^{-7}/\text{afh}$ and the expected cumulative arisings per annum is 10^{-4} then risk band (E) would be selected.

INDIVIDUAL AIRCRAFT PEAK RISK /AFH										
EXPECTED CUMULATIVE ARISINGS / ANNUM										
	10^{-3}	10^{-4}	10^{-5}	10^{-6}	10^{-7}	10^{-8}	10^{-9}	10^{-10}		
	10	1	10^{-1}	10^{-2}	10^{-3}	10^{-4}	10^{-5}	10^{-6}		
	(A)	(B)	(C)	(D)	(E)	(F)	(G)	(H)	(I)	(J)
(1) CIVIL DISASTROUS Uncontrolled flight into highly populated building, crowd or civilian passenger aircraft										
(2) CATASTROPHIC FOR MILITARY TRANSPORT Multiple military or a civilian death or multi-engine aircraft loss										
(3) CRITICAL FOR MILITARY TRANSPORT Jeopardised flight, severe injury, major system damage										
(4) CATASTROPHIC FOR COMBAT AIRCRAFT Loss of pilot or loss of single/twin aircraft or severe injury to civilian passenger										
(5) ATTRITION/CRITICAL FOR COMBAT AIRCRAFT Severe injury, severe occupational illness or major system damage or technically caused loss in combat										
(6) MARGINAL Minor injury, minor occupational illness, or minor system damage										
(7) NEGLIGIBLE Less than minor injury, occupational illness or system damage										

ACCEPTABLE WITHOUT REVIEW

UNDESIRABLE IF MILITARY TRANSPORT

UNDESIRABLE

UNACCEPTABLE

Figure 3: Improved HRI matrix that overcomes the difficulties highlighted by the examples.

4.2 Severity criteria

A problem we have encountered with the severity axis, is with respect to the higher accepted risk of a catastrophe due to technical causes during military active service compared to other flight activities. For example, in a combat situation a pilot might judge it to be necessary to take his aircraft slightly outside the normal flight-control envelope. The issue is that the HRI matrix should identify how much higher a technical risk is then accepted. Assuming a factor of 10 higher risk is accepted, this issue is addressed by broadening the definition of 'critical' (see severity category (5)) to include attrition due to engineering causes (i.e. loss of aircraft during active service), as shown in figure 3.

Secondly, there is a concern that the current definition of the 'catastrophic' category (given in figure 1) is too broad. For example, loss of flight control during a display flight could be far more hazardous than a fuel pump failure of a single engine aircraft (where there is a good chance of being able to glide the aircraft to an unpopulated region before the pilot ejects). There is also public and political pressure to distinguish between the risk of an air accident involving large numbers of civilian casualties and a flight into the ground where the pilot has good chance of being able to eject. To address, this an additional 'civil disastrous' category could be added ('(1)' in figure 3). In addition to hazards associated with flight displays, loss of flight control during take-off, landing or over flight of highly populated areas and buildings could fall into the proposed highest severity category.

Thirdly, there is the option that the two HRI matrices could be subsumed into one table, so that the potential confusion of reading the wrong table is discouraged (as they look very similar). In the proposed matrix, the risk criteria of the current two HRI matrices are preserved so that using the new matrix does not affect the status of risks for which the current matrices work well.

In summary, the proposed developments to the HRI matrix are to use stronger inputs so that it copes better with the wide variety risks that occur in service. The new risk criteria introduced in figure 3, are only intended for suggestion and further discussion. As yet, they are neither implemented in any current analysis nor form part of any operational procedure.

5 OTHER RISK MANAGEMENT CONSIDERATIONS

Application of the HRI matrix is a reactive management approach in so far as it prompts a change in action when there is a change in risk. It is also proactive to the limited extent that risk estimation concerns prediction of the likely exposure to potential hazards. However, another vital proactive aspect of risk management apart from that addressed by the HRI is to minimise of the impact the unforeseeable. A further fundamental issue is that the quality of the statistical methods used can have a very significant impact on the risk values obtained and hence on the HRI category of a risk. These two aspects are discussed in turn.

5.1 Minimisation of the impact of the unforeseeable

Despite conscientious implementations of design standards and service life management procedures, unexpected failures both in the civil and military aircraft fleets have still occurred. Such incidents and particularly in cases where there is a repeat occurrence after a component modification aimed to rectify the problem, can sometimes raise complex questions about whether the level of technical effort invested in lifing such components was adequate. From a technical viewpoint, the complexity of the life limiting behaviour of components combined with the high scatter in their dysfunction life distributions has several implications. Firstly, even using state-of-art analysis there can still be significant error in the estimated service component lives. Secondly, advances over the last decade in: stress analysis standards, lifing methods and revisions of (service usage) exchange rates have brought substantial revisions to a number of military aeroengine component lives. The rapid pace of such progress being made is such that it is likely to continue for the foreseeable future and it is important that fleet managers exploit the resulting benefits.

One of the most important elements of a safe life management procedure for an aeroengine (or indeed for the entire aircraft) is periodic revision of the lives of its critical components. In the immediate term this tends to reduce the severity of costly downward revisions of component lives. In the longer term, experience from the management of large fleets shows that it stops the potential incidence of catastrophic failures rising grossly as the fleet reaches maturity. It also greatly diminishes the likelihood of discovery of cracks in service and subsequent potential grounding of a fleet.

When the reanalysis of a component life has resulted in a significant downward revision, the development of component life extension methods has often played a very significant role in alleviation of the situation (ref. 1). This field concerns enhanced life prediction methods, development of advanced materials, better surface treatments and repair.

Another source of uncertainty is associated with the estimation of the severity of service usage when only a small sample of the fleet is monitored. Recent analyses have shown that much higher (than previously thought) percentage samples of service usage have to be obtained to measure service usage to a certain level of accuracy and thus to support a certain component life.

5.2 Statistical accuracy

As discussed above, estimates of airworthiness risk associated with a particular hazard can be prone to large errors, partly arising from the limited available data upon which a risk model can be based. Lack of data can force assumptions to be made that are highly conservative. However, it normally makes a very significant difference to the quality of risk estimate when it is based on a rigorous statistical analysis, including accounting for sampling error.

To illustrate the sensitivity of some risk estimates, a 3% error in stress, might equate to a 15% error in component life (assuming an SN slope of 5), which in turn might equate to an order of magnitude difference in risk of fatigue failure. If the stress is dependent on a tolerance distribution that has been sampled, then the associated sampling error will have a significant effect on the risk estimate.

6 TRENDS IN MILITARY AIRCRAFT SAFETY AND FUTURE IMPLICATIONS

6.1 Mature aircraft

Using Tornado as an example of an aircraft that has benefited from life management techniques such as those described above the situation is broadly as follows. There has been a fairly large improvement in safety mainly associated with non-technical causes. Although, the improvement in safety associated with technical causes is smaller, this is still a significant technical achievement. Regular revision of aeroengine component lives has resulted in significant changes, and catastrophic failures have been avoided as a result.

For a typical UK military combat aeroengine the estimated total (of the green and amber category) risks might typically be $>3 \times 10^{-5}$ flight hour. That is, up to a factor of 2 greater than the observed rate of technical failures for the whole aircraft. The reason why the estimated total risk of engine failure typically comes out higher than the observed risk of a catastrophic engineering/technical failure of the aircraft is because it is standard practice to calculate risks conservatively.

Based on risk models that QinetiQ has developed, the probability of a fatigue failure per flight hour of a given component occurring at its full life should lie between 2×10^{-8} and 2×10^{-7} . In practice there are typically many complex factors (such as those discussed earlier in the paper) that cause the risks of failure of a few of the components to be sufficiently high to enter the amber or upper green categories of the HRI matrix.

6.2 New aircraft

Essentially the same fatigue life design probabilities are used for EFA (the Eurofighter engine) as for mature aeroengines such as RB199 and Pegasus. The current requirement is that the Eurofighter should achieve a failure rate (due to technical causes) of not more than half of that of a mature aircraft (such as Tornado). This requirement can be met with the current risk criteria in the RAF HRI matrices, if fewer risks are allowed into the amber and upper green HRI categories than is typical for current mature RAF combat aircraft.

6.3 Future aircraft

The relatively low fatalities associated with recent wars involving the RAF and media attention to transport accidents have raised safety expectations. In reference (3), Adelard notes that the UK Health and Safety Executive (HSE) has set military airworthiness targets of reduction of the incidence of fatal and major injury accidents by 5% by 2004 and by 10% by 2010. In the longer term, the stated intent is to

achieve aircrew survival levels of greater than 999 in 1000 per annum by 2050. For future military aircraft, the UK MOD safety document JSP318b (ref. 4) defines the airworthiness criteria as being that:

"The cumulative probability of the loss of an aircraft due to a technical fault and the cumulative probability of a technical failure of the aircraft (inclusive of its systems, structures and stores) which could result in the death of any air crew or its passengers, should both be assessed to be of the order of 1 in a million per flight hour, when operated within the conditions used for the airworthiness system".

The precise magnitude of safety improvement that would be required to meet this criterion should become clear in the fullness of time. Broadly, it indicates that future aircraft design will require the risk of aircraft loss (due to technical causes) to be reduced by in the region of 4 to 20 fold (relative to current mature RAF fighter aircraft). It seems to not be an impossible goal because civil passenger aircraft exceed this safety level. However, military aeroengines are taken frequently to more extreme regions of the flight envelope than a civil aircraft and this can be more severe on the aeroengine components. Achievement of up to a 4 to 20-fold improvement in safety could therefore pose a demanding technical challenge in practice. The lifing standard EJ400 (5) for the EJ200 engine exceeds Def. Stan 00-971 (6) and thus should work to probabilities of component failure that are less than or equal to those specified by lifing standard TU346 for the RB199 engine (7) in Tornado.

High-level safety requirements can only be achieved in practice if their quantitative implications are percolated down to the component-level design standards and HRI procedures. There are two basic options. The option to limit the number of amber category risks to one or two does not seem to be viable, since it could result in grounding of fleets each time more than a couple of risks enter the amber and upper green categories. A more realistic option therefore seems to be to allow revision downwards of the probabilities in the HRI tables according to the required improvement factor over current safety levels. This option would result in proliferation of HRI matrices, which is an additional reason why it is suggested that the two current matrices should be merged into one.

Realisation of a large improvement factor in safety (relative to current mature aircraft) would also require downward revision of the probabilities of failure in Def. Stan. 00-971 (shortly to be absorbed into Def. Stan. 00-970). Based on current risk models, a 20-fold reduction in the risk of fatigue at full life equates to approximately a 15% reduction in component life.

For completeness, a brief mention should be made of the risk management implications for UAV aircraft. Clearly, UAV technology has considerable development potential in terms of flight platforms that: reduce pilot workload, operate more effectively in remote and hazardous environments, allow greater manoeuvrability than a manned aircraft, etc.. Any associated reduction in the need for pilot training could reduce component life requirements and thus alter the optimum lifing tradeoffs between risk, affordability and aeroengine performance.

7 CONCLUSIONS

1. A variety of examples were given to illustrate some of the issues that arise in the risk assessment of aeroengine component failure modes and other potential hazards.
2. The RAF Hazard Risk Index (HRI) management procedure was also critically reviewed.
3. It is suggested that this HRI matrix could be improved using stronger inputs:
 - a. using probabilities of failure of:
 - Peak risk per flight hour (defined in section 3.2.2)
 - Expected cumulative arisings (defined in section 3.2.3)
 - b. and using additional severity criteria that:
 - discriminate risks of technical causes of attrition from other risks;
 - discriminate the risks of more serious catastrophic accidents involving numerous civilian fatalities from those that are limited to flight into the ground and pilot ejection; and
 - enable the two matrices (for fighter aircraft and military transport) to be merged into a unified matrix.

4. A brief review of trends in RAF air safety is given to show why future requirements to improve safety could at some future point necessitate design and risk management standards to be shifted to lower probabilities of failure.

8 REFERENCES

1. A D Boyd-Lee and G F Harrison, 'The development of life extension methods for fracture-critical aero-engine components', in *Qualification of life extension schemes for engine components*, RTO-MP-17, 1999.
2. Gp. Capt. J.K. Meagher and Sqn. Ldr. L. Jones, 'The challenge of managing the logistic support of combat engines in the Royal Air Force', in *Gas turbine operation and technology for land, sea and air propulsion and power systems*, RTO-MP-34, 2000.
3. Validation of Airworthiness Target by the ALARP Principle, Adelard Report Reference D/188/5601/8v1.0, 1 February 2001
4. Regulation of the Airworthiness of Ministry of Defence Aircraft, JSP 318B.
5. Specification for the EJ200-01 R-A1 engine, EJ400, Issue 5, July 1995.
6. UK Military Defence Standards, Def, Stan. 00-971, General Specification for aircraft gas-turbine engines, 1986.
7. Lifting procedures for RB199 group A parts, TU346, issue 3, amendment 4, May 1999.

Paper, 'Risk Assessment Methodologies for Fracture Critical Components' : Discussion

Question from P R Parolo – DGTA, Australia

Given the size of the newly proposed risk/decision matrix, how practical is it to use and how practical is it to quantify a case in so much detail? Based on experiences involving reduced LCF lives in the RAAF, I think that a smaller matrix is more practical for an operator to use.

Presenter's Reply

The matrix in the paper is only a proposal to the RAF.